

An Inductive Proof Method for Simulation-based Compiler Correctness

Sigurd Schneider, Gert Smolka, Sebastian Hack
Saarland Informatics Campus, Saarland University

November 30, 2016

We study induction on the program structure as a proof method for bisimulation-based compiler correctness. We consider a first-order language with mutually recursive function definitions, system calls, and an environment semantics. The proof method relies on a generalization of compatibility of function definition with the bisimulation. We use the inductive method to show correctness of a form of dead code elimination. This is an interesting case study because the transformation removes function, variable, and parameter definitions from the program. While such transformations require modification of the simulation in a coinductive proof, the inductive method deals with them naturally. All our results are formalized in Coq.

1 Introduction

We study induction on the program structure as a proof method for bisimulation-based compiler correctness. We detail inductive equivalence proofs for the language IL [10] with respect to a simple, coinductively defined bisimulation. IL is a first-order language with lexically scoped variables, system calls, mutual recursion, and an environment semantics in the style of Standard ML [6]. The restriction to first-order simplifies the setup of semantics, simulation, and the inductive method. System calls realize (two-way) communication with a system environment and warrant a bisimulation-based notion of program equivalence.

We explain the inductive method by providing paper proofs of the crucial lemma, and a case study that applies the lemma in proofs of (bi)similarities: First, we show contextuality of a bisimulation-based equivalence. Second, we prove correctness of dead code elimination (DCE), which we split into the following two transformations:

1. Unreachable Code Elimination (UCE)
2. Dead Variable Elimination (DVE)

UCE removes unreachable function definitions and unreachable conditional branches. DVE removes dead variable definitions and dead parameters from functions.

We verify correctness with respect to a coinductively defined bisimulation. Intuitively, two configurations are bisimilar if they perform the same (possibly infinite) trace of system calls. Two configurations are similar, if every trace of the left-hand configuration is also a (partial) trace of the right-hand configuration. We show that UCE respects bisimilarity, and that DVE respects simulation.

We use induction on the program structure as a proof method for the correctness arguments. In the inductive proof, the fact that optimizations UCE and DVE remove program statements is no issue. This is in contrast to a coinductive proof method. The coinductive hypothesis cannot be applied without further justification if, for example, the source program reduced, but the target program did not. Such situations naturally arise, for example, if a variable or function definition is removed. The standard solution is index the simulation with a well-founded relation, and allow stutter steps if the well-founded relation is decreased [5]. The main feature of our inductive method is that it does not require modifications of the simulation.

The inductive proof method is enabled by a lemma that strengthens the inductive hypothesis. The plain structural inductive hypothesis cannot readily be used to show the function definition case correct. Suppose \sim is a semantically defined bisimilarity relation for programs. The problem with function definitions is that the following rule of congruence does not directly follow by induction. Function definition introduces a fixed-point of the semantics of s , which requires extra (coinductive) treatment.

$$\frac{s \sim s' \quad t' \sim t'}{\text{fun } f \overline{x} = s \text{ in } t \sim \text{fun } f \overline{x} = s' \text{ in } t'}$$

This rule of congruence (which we prove in Section 7) is a special case of a general lemma we prove, and which we use to strengthen the inductive hypothesis. Our lemma generalizes to optimizations that change function signatures and remove and rename function and variable definitions. In the latter case, the bodies of related functions are generally not equivalent, but their function applications are, provided the arguments are in some other, often asymmetric, relation.

This paper is accompanied by a Coq development which contains formal proofs of all lemmas and theorems. The Coq development is part of a larger compiler verification project and available online:

www.ps.uni-saarland.de/~sdschn/lvc-ind/

1.1 Contributions

The paper makes the following contributions:

1. We develop a simple proof method for IL which supports induction on the program structure for proofs of bisimulation-based equivalence. IL has mutually recursive function definitions and system calls.
2. We detail the method in the proof of contextuality of our simulation, and in the correctness proofs of Unreachable Code Elimination and Dead Variable Elimination. We explain how the method deals with the removal of definitions, function parameters, and mutual recursion.
3. The correctness proofs of UCE and DVE are carried out formally in Coq in a setting with De-Bruijn function binders and mutual recursion. Here in the paper, we present a named version in hope for better readability.

1.2 Related Work

Inductive Proofs for Bisimulations The idea to use compatibility lemmas to simplify correctness proofs is outlined in §2.8 of the master thesis of one of the authors [9]. The master thesis uses a version of IL without mutual recursion and system calls. A basic version of the extension lemma, which enables the inductive method and which we prove in Subsection 6.2, appears in the master thesis as Lemma 3. The master thesis uses the extension lemma to show that contextual equivalence is characterized by a simulation-based definition. In Section 7 of this paper we show that a bisimulation-based definition is sound for contextual equivalence.

Neis et al. [7] recently used an inductive method to deal with stuttering steps when proving their elaborate parametric inter-language simulations (PILS). In the PILS framework, they verify a compiler for an imperative higher-order language with non-mutually recursive functions that take a fixed number of arguments.

Neis et al. use an inductive method to deal with stuttering steps in, among other things, the correctness proof of a form of DCE with respect to PILS that only eliminates unused let-bindings. Neis et al. mention that their framework provides a series of compatibility lemmas simplifying the proof, but do not state the precise form of the lemmas in the paper. Our DCE removes dead function parameters, unused function definitions, and unreachable branches of conditionals.

As Neis et al. deal with a higher-order language in the PILS framework, their setup is necessarily more complicated than ours, and they only give a high-level description of how they setup the induction. This paper aims to explain the inductive method in a simple setting that still allows to see its merit. We include mutual recursion as it directly interacts with the setup of the inductive method. We detail the proofs in the hope to expose the inductive method in general, independent of a framework.

DCE in CompCert Dead code elimination (DCE) in CompCert [4] is carried out by two optimizations. First, the translation from RTL to LTL replaces instructions that write to dead registers with no-ops. Second, the branch tunneling phase removes no-ops.

In CompCert, optimizations that remove instructions are proven correct via a measure argument that justifies applicability of the coinductive hypothesis. Dealing with

a measure is more complicated than a plain coinductive proof. For this reason, only the branch tunneling phase removes instructions. Our inductive approach supports removal of definitions (i.e. instructions) without additional effort in the correctness proof of any optimization and does not require an additional measure.

Correctness Arguments in Verified Compilers The correctness arguments in VeLLVM [12], the verified LLVM project, CompCertSSA [1], and CompCertTSO [11] use exclusively coinduction for correctness proofs. Those compilers operate on a graph-based program representation, so induction on the program structure is not as useful as in our term-based setting.

Howe’s Method Howe’s Method [2] is a general method to show that a (coinductively defined) relation is a congruence. Howe’s method is particularly effective in higher-order settings. Howe’s method first constructs a precongruence candidate relation that contains bisimilarity and can easily be shown to be a congruence. Afterwards the candidate relation is shown to coincide with bisimilarity. In our work we prove by coinduction that for showing two function definitions bisimilar it suffices to show that their bodies are bisimilar, assuming that all related functions in the environment are bisimilar. Howe’s method seems to be geared towards congruence properties and we are not aware of work extending it to optimizations that change function signatures.

CakeML CakeML [8] is a verified compiler for a substantial subset of Standard ML. CakeML originally uses big-step semantics, which does not account for diverging behaviors. Recently, CakeML switched to an evaluation function with a step limit to specify the semantics. Both approaches directly support inductive proofs on the semantics.

1.3 Outline

The paper is organized as follows. We define the syntax and semantics of the language IL in Section 2. In Section 3 we repeat the definition of program equivalence from previous work [10], and give a new characterization using parameterized co-induction [3]. We then prove compatibility rules admissible that we use repeatedly in the following proofs. We develop the inductive method in Section 6 and use it in Section 7 to show that the bisimulation we defined is contextual. We describe how program analysis information is represented in our framework in Subsection 8.1. We specify reachability and prove unreachable code elimination correct in Section 8. We specify true liveness and prove dead variable elimination correct in Section 9. We discuss the formal development in Section 10 and conclude in Section 11.

2 IL

2.1 Values, Variables, and Expressions

We assume a type \mathbb{V} of values and a function $\beta : \mathbb{V} \rightarrow \mathbb{B} = \{\mathbf{true}, \mathbf{false}\}$ that we use to simplify the semantic rule for the conditional. By convention, v ranges over \mathbb{V} . We use the countably-infinite alphabet \mathcal{V} for names x, y, z of values, which we call *variables*.

$\eta ::= e \mid \alpha(\bar{e})$	extended expression
$\text{Term} \ni s, t ::= \text{let } x = \eta \text{ in } s$	variable binding
$\mid \text{if } e \text{ then } s \text{ else } t$	conditional
$\mid e$	expression
$\mid \text{fun } \overline{f \bar{x} = s} \text{ in } t$	function definition
$\mid f \bar{e}$	application

Figure 1: Syntax of IL

We assume a type Exp of expressions. By convention, e ranges over Exp . Expressions are pure, their evaluation is deterministic and may fail, hence expression evaluation is a function $\llbracket \cdot \rrbracket : \text{Exp} \rightarrow (\mathcal{V} \rightarrow \mathbb{V}_\perp) \rightarrow \mathbb{V}_\perp$. Environments are of type $\mathcal{V} \rightarrow \mathbb{V}_\perp$ to track uninitialized variables, and are partially ordered by \sqsubseteq , which is the pointwise lifting of the relation defined by the two equations $\perp \sqsubseteq w$ and $w \sqsubseteq w$, where $w \in \mathbb{V}_\perp$. We assume that expression evaluation is monotone, i.e., $V \sqsubseteq V' \rightarrow \llbracket e \rrbracket V \sqsubseteq \llbracket e \rrbracket V'$.

We assume a function $\text{fv} : \text{Exp} \rightarrow \text{set } \mathcal{V}$ such that for all environments V, V' that agree on $\text{fv}(e)$ we have $\llbracket e \rrbracket V = \llbracket e \rrbracket V'$. We use the notation \bar{x} for a list of variables. We lift $\llbracket \cdot \rrbracket$ pointwise to lists of expressions in a strict fashion: $\llbracket \bar{e} \rrbracket$ yields a list of values if none of the expressions in \bar{e} failed to evaluate, and \perp otherwise.

We sometimes omit the side condition $\llbracket \bar{e} \rrbracket V \neq \perp$ in the presentation if $\llbracket \bar{e} \rrbracket V$ is used in a place where type \mathbb{V} is required. For example, we write $\beta(\llbracket \bar{e} \rrbracket V) = \mathbf{true}$ instead of $\exists v : \mathbb{V}, \llbracket \bar{e} \rrbracket V = v \wedge \beta v = \mathbf{true}$.

2.2 Syntax

IL is a first-order language with a tail-call restriction, mutual recursion, and system calls. IL syntactically enforces a first-order discipline by using a separate alphabet \mathcal{F} for function names f, g, h . Variables are lexically scoped binders, and a function definition creates a closure that captures variables.

IL uses a third alphabet \mathcal{A} for names α which we call *actions*. The term $\text{let } x = \alpha(\bar{e}) \text{ in } \dots$ is like a system call α with argument list \bar{e} that non-deterministically returns a value.

IL allows mutually recursive function definitions. The syntax of IL is given in Figure 1.

2.3 Semantics

The semantics of IL is given as small-step relation \longrightarrow in Figure 2. Note that the tail-call restriction ensures that no call stack is required. The reduction relation \longrightarrow operates on **configurations** of the form (L, V, s) where s is the IL term to be evaluated. The semantics does not rely on substitution, but uses an environment $V : \mathcal{V} \rightarrow \mathbb{V}_\perp$.

$$\begin{array}{c}
\text{OP} \\
\hline
\frac{\llbracket e \rrbracket V = v}{L \mid V \mid \text{let } x = e \text{ in } s \longrightarrow L \mid V[x \mapsto v] \mid s} \\
\\
\text{COND} \\
\hline
\frac{\llbracket e \rrbracket V = v \quad \beta(v) = b}{L \mid V \mid \text{if } e \text{ then } s_{\text{true}} \text{ else } s_{\text{false}} \longrightarrow L \mid V \mid s_i} \\
\\
\text{EXTERN} \\
\hline
\frac{v' \in \mathbb{V} \quad \llbracket \bar{e} \rrbracket V = \bar{v}}{L \mid V \mid \text{let } x = \alpha(\bar{e}) \text{ in } s \xrightarrow{v' = \alpha(\bar{v})} L \mid V[x \mapsto v'] \mid s} \\
\\
\text{FUN} \\
\hline
\frac{}{L \mid V \mid \text{fun } \overline{f \bar{x} = s} \text{ in } t \longrightarrow (\overline{f \bar{x} = s})_V; L \mid V \mid t} \\
\\
\text{APP} \\
\hline
\frac{\llbracket \bar{e} \rrbracket V = \bar{v} \quad L_f = (V', \bar{x}, s)}{L \mid V \mid f \bar{e} \longrightarrow L^{-f} \mid V'[\bar{x} \mapsto \bar{v}] \mid s} \\
\\
\hline
(\overline{f \bar{x} = s})_V = [\overline{f : (V, \bar{x}, s)}]
\end{array}$$

Figure 2: Semantics of IL

for variable definitions and a context L of function definitions. Transitions in \longrightarrow are labeled with *events* ϕ . By convention, ψ ranges over events different from τ .

$$\mathcal{E} \ni \phi ::= \tau \mid v = \alpha(\bar{v})$$

The silent event is denoted by τ , and we omit it by convention.

A **context** is a list of groups of named definitions. For example, the context $K = [f_1 : a_1, f_2 : a_2]; [g_1 : b_1]$ consists of three definitions in two groups. We define a function *dom* that yields the domain of a context as a list, e.g. $\text{dom } K = f_1, f_2, g_1$. A definition in a context may refer to previous definitions and definitions in its group. Notationally, we use contexts like functions: To access the first element with name f , we write $\boxed{L_f}$ and we have $L_f = \perp$ if no such element exists. We write L^{-f} for the context obtained from L by dropping all groups before the first group containing f . We write $;$ for context concatenation and \emptyset for the empty context.

A **closure** is a tuple $(V, \bar{x}, s) \in \mathcal{C}$ consisting of an environment V , a parameter list \bar{x} , and a function body s . Since a function f in a context can only refer to previously defined functions and functions in its own group, the first-order restriction allows the closures to be non-recursive: function closures do not need to close under functions. An application $f\bar{e}$ causes the function context L to rewind to L^{-f} , i.e. up to the group with the definition of f (rule APP). In contrast to higher-order formulations, we do not define closures mutually recursively with the values of the language.

A **system call** $\text{let } x = \alpha \bar{e} \text{ in } s$ invokes a function α of the system, which is not assumed to be deterministic. This reflects in the rule EXTERN, which does not restrict the result value of the system call other than requiring that it is a value. The transition records the system call name α , the argument values \bar{v} and the result value v' in the event $v' = \alpha(\bar{v})$.

3 Program Equivalence

Before any transformation can be proven correct, we must formally define what semantic equivalence means. Semantic equivalence is not directly tight to the language, but only to the way the language interacts with its environment. In our case, the language interacts with the environment via system calls, and possibly a result value. We abstract this behavior with internally deterministic reduction systems (IDRS), that we previously introduced [10].

Definition 1 A *reduction system* (RS) is given by a tuple $(\Sigma, \mathcal{E}, \longrightarrow, \tau, \text{res})$ such that

1. $(\Sigma, \mathcal{E}, \longrightarrow)$ is a LTS
2. $\text{res} : \Sigma \rightarrow \mathbb{V}_\perp$
3. $\text{res } \sigma = v \Rightarrow \sigma \not\rightarrow$
4. $\tau \in \mathcal{E}$

An *internally deterministic* reduction system (IDRS) additionally satisfies

5. $\sigma \xrightarrow{\phi} \sigma_1 \wedge \sigma \xrightarrow{\phi} \sigma_2 \Rightarrow \sigma_1 = \sigma_2$ action-deterministic

BISIM-SILENT			BISIM-TERM	
$\sigma_1 \longrightarrow^+ \sigma'_1$	$\sigma_2 \longrightarrow^+ \sigma'_2$	$\sigma'_1 \sim \sigma'_2$	$\sigma_1 \Downarrow w$	$\sigma_2 \Downarrow w$
$\sigma_1 \sim \sigma_2$			$\sigma_1 \sim \sigma_2$	
BISIM-EXTERN			SIM-ERROR	
$\sigma_1 \longrightarrow^* \sigma'_1$		$\sigma'_1 \rightsquigarrow \sigma'_2$		$\sigma'_1 \text{ terminal}$
$\sigma_2 \longrightarrow^* \sigma'_2$	$\sigma'_1, \sigma'_2 \text{ ready}$	$\sigma'_2 \rightsquigarrow \sigma'_1$	$\sigma_1 \longrightarrow^* \sigma'_1$	$\text{res } \sigma'_1 = \perp$
$\sigma_1 \sim \sigma_2$			$\sigma_1 \lesssim \sigma_2$	

Figure 3: Defining Rules of Similarity and Bisimilarity

$$6. \sigma \xrightarrow{\phi} \sigma_1 \wedge \sigma \xrightarrow{\tau} \sigma_2 \Rightarrow \phi = \tau \quad \tau\text{-deterministic}$$

The semantics of IL forms an IDRS: We define *res* such that $\text{res}(\sigma) = v$ if σ is of the form (F, V, e) and $\llbracket e \rrbracket V = v$. Otherwise, $\text{res}(\sigma) = \perp$.

3.1 Similarity and Bisimilarity

To define what it means that two IDRS behave equivalently, we use (bi)similarity. Bisimilarity is obtained as the greatest fixed-point, and naturally accounts for diverging behaviors. In previous work [10] we have given a definition of the bisimilarity relation we present here, and we have shown that it sound and complete for trace equivalence.

Before we give the rules defining (bi)similarity, need some definitions. We write $\sigma \Downarrow w$ (where $w \in \mathbb{V}_\perp$) if σ terminates with w , that is, $\sigma \longrightarrow^* \sigma'$ such that σ' is \longrightarrow -terminal and $\text{res}(\sigma') = w$. We also want to be able to identify configurations which are about to execute a system call, and say that such configurations are *ready*. Finally we introduce the notation $\sigma_1 \overset{R}{\rightsquigarrow} \sigma_2$ for the standard forward-simulation property. That is, every transition σ_1 takes can also be taken by σ_2 , and the two successor configurations are related by R .

Definition 2 (Bisimilarity) Let $(S, \mathcal{E}, \longrightarrow, \text{res}, \tau)$ be an IDRS. We define bisimilarity in type theory as relation of type $S \rightarrow S \rightarrow \mathbf{P}$, where \mathbf{P} is the universe of propositions. Bisimilarity \sim is defined coinductively as the greatest relation closed under the rules BISIM-SILENT, BISIM-EXTERN, BISIM-TERM in Figure 3.

BISIM-SILENT allows to match finitely many steps on both sides, as long as all transitions are silent. This makes sense for IDRS, but would not yield a meaningful definition otherwise. BISIM-EXTERN ensures that every external transition of σ'_1 is matched by the same external transition of σ'_2 , and vice versa. This ensures that if two programs are in relation, they react to every possible result value of the external call in a bisimilar way. σ'_1, σ'_2 are required to be ready to simplify case distinctions by ensuring that the next event cannot be τ .

Definition 3 Let $(S, \mathcal{E}, \longrightarrow, res, \tau)$ be an IDRS. Similarity is defined as the greatest relation closed under the rules BISIM-SILENT, BISIM-EXTERN, BISIM-TERM and SIM-ERROR in Figure 3.

SIM-ERROR can be used to justify similarity for any configuration on the right side, if the left side can be shown to reduce to a stuck configuration.

3.2 Bisimilarity as Symmetrization of Similarity

Obtaining bisimilarity as symmetrization of similarity is useful if the properties one wants to show are symmetric properties: Bisimilarity is obtained from a proof of similarity and symmetry. If the property is not symmetric, one needs two proofs of similarity, which, in practice, share a lot of arguments. Leroy [4] and Sevcík [11] avoid a second proof for the backward direction by showing that on the class of LTS they are using, forward and backward simulation coincide. In our setting, bisimilarity is the basic definition, and simulation is obtained by adding an “escape” rule (SIM-ERROR) that justifies similarity if the left configuration is stuck. In this way, we can show forward and backward direction in one proof, but do not require the two directions to be equivalent. In the presence of non-determinism, the two forward and backward simulation do not coincide.

4 Parameterized Coinduction

For the formalization, we need to define simulation and bisimulation via parameterized coinduction [3] to side-step the too restrictive guardedness check for co-fixed points in Coq. A coinductively defined function must be productive to be well-formed, a criterion that is dual to the requirement that an inductively defined function must be terminating. Coq requires co-recursion to occur syntactically directly below a constructor of the co-inductive definition, which is a sufficient criterion for productivity. Parameterized coinduction allows for productivity to be accounted for in a semantic way. We recapitulate the basic setup of parameterized coinduction following Hur et al. [3] in this section, and outline how parameterized coinduction works in Remark 1, when we have all definitions at hand.

Definition 4 (Complete Prelattice) A complete prelatice $(X, \sqsubseteq, \sqcap, \sqcup, \top, \perp)$ is a complete lattice that is defined with respect to $x \equiv y := x \sqsubseteq y \wedge y \sqsubseteq x$ instead of equality, i.e. a lattice which does not require anti-symmetry.

The setup relies on the notion of a complete prelatice. Hur et al. do not require anti-symmetry, but base the paper presentation on a complete lattice nonetheless. We apply parameterized coinduction to functions into \mathbf{P} , the universe of propositions. Function types into \mathbf{P} only form a complete lattice, if the axioms of propositional extensionality and functional extensionality are assumed. The function types into \mathbf{P} each form a complete prelatice.

Definition 5 (Greatest Fixed Point) Let X be a complete prelatattice. We define a function

$$\begin{aligned}\text{cofix} &: (X \rightarrow X) \rightarrow X \\ \text{cofix } f &:= \bigsqcup \{y \in X \mid y \sqsubseteq fy\}\end{aligned}$$

We use the notations $\nu x.s := \text{cofix}(\lambda x.s)$ and $\nu f := \text{cofix } f$.

Fact 1 Let X be a complete prelatattice and f be a monotone function. Then $\text{cofix } f \sqsubseteq f(\text{cofix } f)$.

Definition 6 (Parameterized Greatest Fixed Point) Let X be a complete prelatattice and $f : X \rightarrow X$ be a monotone function. We define a function

$$\begin{aligned}\mathbf{G} &: (X \xrightarrow{\text{mon}} X) \xrightarrow{\text{mon}} X \xrightarrow{\text{mon}} X \\ \mathbf{G} f x &:= \nu y.f(x \sqcup y)\end{aligned}$$

It is easy to check that \mathbf{G} and $\mathbf{G}f$ are monotone.

Lemma 1 (Initialize) $\nu f \equiv \mathbf{G}f \perp$.

Lemma 2 (Unfold) $\mathbf{G}fx \equiv f(x \sqcup \mathbf{G}fx)$.

Lemma 3 (Accumulate) $y \sqsubseteq \mathbf{G}fx \leftrightarrow y \sqsubseteq \mathbf{G}f(x \sqcup y)$.

Proof. See [3]. □

Corollary 1 If $\forall z, x \sqsubseteq z \rightarrow y \sqsubseteq z \rightarrow y \sqsubseteq \mathbf{G}fz$ then $y \sqsubseteq \mathbf{G}fx$.

Remark 1 outlines the usage of Corollary 1 as coinductive proof principle. The definition of \mathbf{G} and its lemmas are provided by the Paco library [3]. The Paco library realizes \mathbf{G} directly as a coinductively defined predicate, instead of using the cofixed point operator we defined for this presentation in Definition 5.

5 Similarity and Bisimilarity as Parameterized Greatest Fixed Point

We obtain definitions equivalent to similarity and bisimilarity with the fixed point operator \mathbf{G} from a single function. The use of a single function allows us to show many properties which hold for both, similarity and bisimilarity, with one lemma. This saves a lot of repetition particularly in the proof of transitivity.

Definition 7 We define the function *sim* that generates similarity and bisimilarity in Figure 4.

$\text{STy} \ni s ::= \text{bisim} \mid \text{sim}$	
$\text{sim} : (\text{STy} \rightarrow \Sigma \rightarrow \Sigma \rightarrow \mathbf{P}) \rightarrow (\text{STy} \rightarrow \Sigma \rightarrow \Sigma \rightarrow \mathbf{P})$	
$\text{sim } r p \sigma_1 \sigma_2 := (\exists w. \sigma_1 \Downarrow w \wedge \sigma_2 \Downarrow w)$	BISIM-TERM
$\vee (\exists \sigma'_1 \sigma'_2. \sigma_1 \longrightarrow^+ \sigma'_1 \wedge \sigma_2 \longrightarrow^+ \sigma'_2 \wedge r p \sigma'_1 \sigma'_2)$	BISIM-STEP
$\vee (\exists \sigma'_1 \sigma'_2. \sigma_1 \longrightarrow^+ \sigma'_1 \wedge \sigma_2 \longrightarrow^+ \sigma'_2$	
$\wedge \sigma'_1, \sigma'_2 \text{ ready} \wedge \sigma'_1 \overset{r p}{\rightsquigarrow} \sigma'_2 \wedge \sigma'_2 \overset{r p}{\rightsquigarrow} \sigma'_1)$	BISIM-EXTERN
$\vee (p = \text{sim}$	
$\wedge \exists \sigma'_1. \sigma_1 \longrightarrow^* \sigma'_1 \wedge \sigma'_1 \text{ terminal} \wedge \text{res } \sigma'_1 = \perp)$	SIM-ERROR

Figure 4: Generating Function for Simulation and Bisimulation. Each disjunct corresponds to a rule from Figure 3.

Remark 1 (Outline of Parametric Co-Induction) A parameterized coinduction using sim always has a conclusion of the form $R \subseteq \mathbf{G} \text{sim } r p$ for some relations R and r . Applying Corollary 1 sets up the coinduction: We have to show $R \subseteq \mathbf{G} \text{sim } r' p$ but can assume $r \subseteq r'$ and $R \subseteq r'$. The assumption $R \subseteq r'$ is the coinductive hypothesis. The proof typically proceeds by unfolding \mathbf{G} according to Lemma 2: $R \subseteq \text{sim}(r' \cup \mathbf{G} \text{sim } r') p$. Unfolding exposes the generating function sim , each disjunct of which corresponds to a constructor (cf. Figure 3). In places where the constructor uses co-recursion, the function sim applies its parameter r . In our proof, the parameter is $r' \cup \mathbf{G} \text{sim } r'$. This ensures that the co-hypothesis $R \subseteq r'$ is only applied after one of the constructors has been “used”. The parameter in the definition of \mathbf{G} encodes the productivity requirement semantically.

Lemma 4 Let $p : \text{STy}$ and $\sigma_1, \sigma_2, \sigma_3 : \Sigma$. If $\mathbf{G} \text{sim} \perp p \sigma_1 \sigma_2$ and $\mathbf{G} \text{sim} \perp p \sigma'_2 \sigma_3$ and $\sigma_2 \longrightarrow^* \sigma'_2$ or $\sigma'_2 \longrightarrow^* \sigma_2$ then $\mathbf{G} \text{sim} \perp p \sigma_1 \sigma_3$.

Proof. The proof is by case analysis on $\mathbf{G} \text{sim} \perp p \sigma_1 \sigma_2$ and $\mathbf{G} \text{sim} \perp p \sigma'_2 \sigma_3$. The cases are not difficult, but tedious. \square

Definition 8 Let $r : \text{STy} \rightarrow \Sigma \rightarrow \Sigma \rightarrow \mathbf{P}$. We define:

$$\begin{aligned} \approx_r^p &:= \mathbf{G} \text{sim } r p \\ \lesssim_r &:= \approx_r^{\text{sim}} \\ \sim_r &:= \approx_r^{\text{bisim}} \end{aligned}$$

Lemma 5 \lesssim_\perp is a preorder.

Proof. Reflexivity is trivial; transitivity is an instance of Lemma 4. \square

Lemma 6 \sim_\perp is an equivalence relation.

SIM-EXPANSION-CLOSED			SIM-RETRACT		
$\sigma_1 \longrightarrow^* \sigma'_1$	$\sigma_2 \longrightarrow^* \sigma'_2$	$\sigma'_1 \approx_r^p \sigma'_2$	$\sigma_1 \longrightarrow \sigma''_1$	$\sigma_2 \longrightarrow \sigma''_2$	
$\sigma_1 \approx_r^p \sigma_2$			$\sigma'_1 \longrightarrow \sigma''_1$	$\sigma'_2 \longrightarrow \sigma''_2$	$\sigma_1 \approx_r^p \sigma_2$
			$\sigma'_1 \approx_r^p \sigma'_2$		

Figure 5: Closedness under Expansion and Retraction

Proof. Reflexivity and symmetry are trivial; transitivity is an instance of Lemma 4. \square

The following theorem establishes trust in our non-standard setup. The definitions obtained from parameterized coinduction and the function *sim* are equivalent to the more basic definitions from Subsection 3.1.

Lemma 7 $\lesssim \equiv \lesssim_\perp$ and $\sim \equiv \sim_\perp$.

5.1 Properties of Similarity and Bisimilarity

The following admissible rules allow us to retract to reduction successors of states when showing (bi)similarity.

Lemma 8 The rules in Figure 5 are admissible.

5.2 Structural Rules

The following lemmas are formulated with respect to \approx_r^p , which allows us to use one proof to show a property of both simulation and bisimulation.

Lemma 9 The rules in Figure 6 are admissible.

Proof. We only show SIM-LET-OP. After rewriting with Lemma 2, we have to show that $(L, V, \text{let } x = e \text{ in } s)$ and $(L', V', \text{let } x' = e' \text{ in } s')$ are related by $\text{sim}(r \cup \approx_r^p)p$. Case analysis on $\llbracket e \rrbracket V$.

- Case $\llbracket e \rrbracket V = v$. We unfold *sim* and show the case BISIM-SILENT. The two required successor states exist:
 1. $(L, V, \text{let } x = e \text{ in } s) \longrightarrow^+ (L, V[x \mapsto v], s)$
 2. $(L', V', \text{let } x' = e' \text{ in } s') \longrightarrow^+ (L', V'[x' \mapsto v], s')$ $(L, V[x \mapsto v], s) (\approx_r^p \cup r) (L', V'[x' \mapsto v], s')$ holds by assumption, which finishes the case.
- Case $\llbracket e \rrbracket V = \perp$. We unfold *sim* and show the case BISIM-TERM. Both states are terminal, and the way we defined the result function ensures that $(L, V, \text{let } x = e \text{ in } s) \Downarrow \perp$ and $(L', V', \text{let } x' = e' \text{ in } s') \Downarrow \perp$. \square

$$\begin{array}{c}
\text{SIM-LET-OP} \\
\frac{\llbracket e \rrbracket V = \llbracket e' \rrbracket V' \quad \forall v, (L, V[x \mapsto v], s) (\approx_r^p \cup r) (L', V'[x' \mapsto v], s')}{(L, V, \text{let } x = e \text{ in } s) \approx_r^p (L', V', \text{let } x' = e' \text{ in } s')} \\
\\
\text{SIM-LET-CALL} \\
\frac{\llbracket \bar{e} \rrbracket V = \llbracket \bar{e}' \rrbracket V' \quad \forall v, (L, V[x \mapsto v], s) (\approx_r^p \cup r) (L', V'[x' \mapsto v], s')}{(L, V, \text{let } x = f \bar{e} \text{ in } s) \approx_r^p (L', V', \text{let } x' = f \bar{e}' \text{ in } s')} \\
\\
\text{SIM-COND} \\
\frac{\llbracket e \rrbracket V = \llbracket e' \rrbracket V' \quad \beta(\llbracket e \rrbracket V) = \mathbf{true} \rightarrow (L, V, s) (\approx_r^p \cup r) (L', V', s') \quad \beta(\llbracket e \rrbracket V) = \mathbf{false} \rightarrow (L, V, t) (\approx_r^p \cup r) (L', V', t')}{(L, V, \text{if } e \text{ then } s \text{ else } t) \approx_r^p (L', V', \text{if } e' \text{ then } s' \text{ else } t')}
\end{array}$$

Figure 6: Admissible Rules

We prove in general that conditionals can be eliminated if the value of the condition is statically known.

Lemma 10 If

- $\beta(\llbracket e \rrbracket \emptyset) = \perp \rightarrow \llbracket e \rrbracket V = \llbracket e \rrbracket V'$
- $\forall v, \llbracket e \rrbracket V = v \rightarrow \beta v = \mathbf{true} \rightarrow \beta(\llbracket e \rrbracket \emptyset) \neq \mathbf{false} \rightarrow (L, V, s_1) \approx_r^p (L', V', s'_1)$
- $\forall v, \llbracket e \rrbracket V = v \rightarrow \beta v = \mathbf{false} \rightarrow \beta(\llbracket e \rrbracket \emptyset) \neq \mathbf{true} \rightarrow (L, V, s_2) \approx_r^p (L', V', s'_2)$

then

$$\begin{array}{c}
(L, V, \quad \text{if } e \text{ then } s_1 \text{ else } s_2) \\
\approx_r^p \quad (L', V', \quad \text{if } \llbracket e \rrbracket \emptyset = \mathbf{true} \text{ then } s'_1 \\
\quad \quad \quad \text{else if } \llbracket e \rrbracket \emptyset = \mathbf{false} \text{ then } s'_2 \\
\quad \quad \quad \text{else if } e \text{ then } s'_1 \text{ else } s'_2).
\end{array}$$

Proof. Case analysis on $\beta(\llbracket e \rrbracket \emptyset)$.

- Case $\beta(\llbracket e \rrbracket \emptyset) = \mathbf{true}$. By monotonicity of expression evaluation, there is v such that $\llbracket e \rrbracket V = v$ and $\beta v = \mathbf{true}$. We apply SIM-EXPANSION-CLOSED, reducing only the right side one step and finish with the second assumption.
- Case $\beta(\llbracket e \rrbracket \emptyset) = \mathbf{false}$. Analogous to the previous case.
- Case $\beta(\llbracket e \rrbracket \emptyset) = \perp$. Case analysis on $\llbracket e \rrbracket V$. If $\llbracket e \rrbracket V = \perp$, both sides are stuck by the first assumption. We unfold via Lemma 2 and use the case SIM-TERM of *sim* to show simulation. If $\llbracket e \rrbracket V = v$, then $\llbracket e \rrbracket V' = v$ by the first assumption. Case analysis on βv . If $\beta v = \mathbf{true}$ ($\beta v = \mathbf{false}$) we apply SIM-EXPANSION-CLOSED to reduce both sides one step and finish with the second (third) assumption. \square

<pre> fun f (x,y) = if (x > 9) then 1 else f (x+1, y) in f (3,2) </pre>	<pre> fun f (x) = if (x > 9) then 1 else f (x+1) in f (3) </pre>
--	---

Figure 7: An example program before (left) and after (right) dead variable elimination.

6 An inductive proof method

We develop an inductive proof method for (bi)similarity, i.e. for statements of the form $(L, V, s) \approx_r^p (L', V', s')$. Such proofs require relating function contexts L, L' , and we will use the relation $L \ r \ L' :^{\mathcal{P}} \Lambda$ defined below. For flexibility, the relation is parameterized by a so called *proof relation* \mathcal{P} , which describes which functions in L, L' are related, and how their arguments and parameters differ.

Definition 9 (Proof Relation) A proof relation is a tuple $(A, Param, Arg, Idx)$ such that

1. $A : \text{Type}$
2. $Param : A \rightarrow \overline{V} \rightarrow \overline{V} \rightarrow \mathbf{P}$
3. $Arg : A \rightarrow \overline{V} \rightarrow \overline{V} \rightarrow \mathbf{P}$
4. $Idx : A \rightarrow \mathcal{F} \rightarrow \mathcal{F} \rightarrow \mathbf{P}$

The proof relation is indexed by a type A , which typically represents program analysis information. A proof relation defines conditions on formal parameters ($Param$), arguments at function calls (Arg), and function names (Idx). For example, consider Figure 7, which contains a program before (left) and after dead variable elimination (right), an optimization we verify in Section 9. The proof relation in Definition 20 relates the two versions of the function f and expresses that, for example, y is removed from the function parameters of f because the second parameter is not used in the body of f .

6.1 Relating Function Contexts

We define the relation $L \ r \ L' :^{\mathcal{P}} \Lambda$, which intuitively means that if two functions from L, L' are related by Idx , then their parameters satisfy the relation $Param$ and the functions are equivalent if called with arguments in relation Arg .

Definition 10 Given a proof relation \mathcal{P} and analysis information context Λ , and function context L, L' we say Λ and L, L' are in parameter relation with respect to \mathcal{P} , written $\boxed{Param \ \Lambda \ L \ L'}$, if whenever $Idx \ \Lambda \ f \ f'$ and $L_f = (V, \overline{x}, s)$ and $L'_{f'} = (V', \overline{x}', s')$ then $Param \ \Lambda_f \ \overline{x} \ \overline{x}'$.

$Param \ \Lambda \ L \ L'$

$\text{App}^{\mathcal{P}} \Lambda L L'$

Definition 11 Given a proof relation \mathcal{P} , and function context L, L' , and analysis information context Λ , we define a relation $\boxed{\text{App}^{\mathcal{P}} \Lambda L L'}$ on configurations such that

$$\begin{aligned} & \text{App}^{\mathcal{P}} \Lambda L L' (L, V, f \bar{e}) (L', V', f' \bar{e}') \\ & :\leftrightarrow \exists V' \bar{x} \bar{x}' s', L_f = (V, \bar{x}, s) \wedge L'_{f'} = (V', \bar{x}', s') \\ & \wedge \text{Idx} \Lambda_f f f' \wedge \text{Arg} \Lambda_f (\llbracket \bar{e} \rrbracket V) (\llbracket \bar{e}' \rrbracket V') \\ & \wedge |\bar{x}| = |\bar{e}| \wedge |\bar{x}'| = |\bar{e}'| \end{aligned}$$

The relation $\text{App}^{\mathcal{P}} \Lambda L L'$ relates application configurations that satisfy the requirements imposed by the proof relation.

$L r L' :^{\mathcal{P}} \Lambda$

Definition 12 Two function contexts L, L' are in r -relation with respect to Λ and \mathcal{P} , written $\boxed{L r L' :^{\mathcal{P}} \Lambda}$ if

1. $\text{dom} \Lambda = \text{dom} L$
2. $\text{Param} \Lambda L L'$
3. $\text{Idx} \Lambda_f f f' \rightarrow (f \in \text{dom} L \leftrightarrow f' \in \text{dom} L')$
4. $\text{App}^{\mathcal{P}} \Lambda L L' \subseteq r$

Lemma 11 If $r \subseteq r'$ and $L r L' :^{\mathcal{P}} \Lambda$ then $L r' L' :^{\mathcal{P}} \Lambda$.

6.2 Extending Related Function Contexts

We prove the central lemma of our inductive proof method, which we call the *extension lemma*. When descending under function definitions, related L and L' are extended with new closures. The inductive hypothesis provides that the bodies of these functions are (bi)similar, but this does not readily mean that the corresponding semantic fixed-points are (bi)similar. The extension lemma (Lemma 15) accounts for the semantics of the fixed-point operator.

$\text{Bdy}_{L,L'}^{\mathcal{P}} \Lambda K K'$

Definition 13 Given a proof relation \mathcal{P} , function context L, L' and K, K' , and analysis information Λ , we define a relation $\boxed{\text{Bdy}_{L,L'}^{\mathcal{P}} \Lambda K K'}$ on configurations such that

$$\begin{aligned} & \text{Bdy}_{L,L'}^{\mathcal{P}} \Lambda K K' (K; L, V[\bar{x} \mapsto v], s) (K; L', V'[\bar{x}' \mapsto v'], s') \\ & :\leftrightarrow \exists f f' V V', K_f = (V, \bar{x}, s) \wedge K'_{f'} = (V', \bar{x}', s') \\ & \wedge \text{Idx} \Lambda_f f f' \wedge \text{Arg} \Lambda_f \bar{v} \bar{v}' \end{aligned}$$

The relation $\text{Bdy}_{K,K'}^{\mathcal{P}} \bar{a} F F'$ relates configurations that are obtained by one reduction from application configurations that satisfy the requirements imposed by the proof relation. We set up our inductive proofs such that the inductive hypothesis provides that these configurations are equivalent.

$$K \parallel \Lambda; \Lambda' \parallel_{\mathcal{P}} K'$$

Definition 14 A proof relation \mathcal{P} separates two contexts K, K' under $\Lambda; \Lambda'$, written

$$\boxed{K \parallel \Lambda; \Lambda' \parallel_{\mathcal{P}} K'} \text{ if:}$$

1. $\text{dom } \Lambda = \text{dom } K$
2. $\text{Idx } (\Lambda; \Lambda')_f f f' \rightarrow (f \in \text{dom } K \leftrightarrow f' \in \text{dom } K')$

Lemma 12 (Extending Parameter Relations) If we have $\text{Param } \Lambda K K'$ and also $K \parallel \Lambda; \Lambda' \parallel_{\mathcal{P}} K'$ and $\text{Param } \Lambda' L L'$ then $\text{Param } (\Lambda; \Lambda') (K; L) (K'; L')$.

Separation requires that functions in K are only related to functions in K' , and vice versa.

Lemma 13 Let \mathcal{P} be a proof relation. If 1. $K \parallel \Lambda; \Lambda' \parallel_{\mathcal{P}} K'$, 2. $\text{Param } \Lambda K K'$, 3. $\text{Bdy}_{L, L'}^{\mathcal{P}} (\Lambda; \Lambda') K K \subseteq (\approx_r^{\mathcal{P}} \cup r)$ and 4. $L \approx_r^{\mathcal{P}} L' :^{\mathcal{P}} \Lambda'$ then $K; L \approx_r^{\mathcal{P}} K; L' :^{\mathcal{P}} \Lambda; \Lambda'$.

Proof. The proof distinguishes whether the function pair is from K and K' or L and L' . This is possible because \mathcal{P} separates K, K' under Λ . In the first case, the result follows from (3) after a lock-step simulation step that reduces function applications on both sides. In the second case, the result follows from (4) and SIM-RETRACT. \square

Definition 15 Given a proof relation \mathcal{P} , function context K, K' are in r -relation under L and L' with respect to \mathcal{P} and $\Lambda; \Lambda'$, written $\boxed{L \mid L' \vdash K r K' :^{\mathcal{P}} \Lambda; \Lambda'}$ if

1. $K \parallel \Lambda; \Lambda' \parallel_{\mathcal{P}} K'$
2. $\text{Param } \Lambda K K'$
3. $\forall r, K; L r K'; L' :^{\mathcal{P}} \Lambda; \Lambda' \rightarrow \text{Bdy}_{L, L'}^{\mathcal{P}} K K' (\Lambda; \Lambda') \subseteq r$

Lemma 14 (Fix Compatibility) Let \mathcal{P} be a proof relation. If $L \mid L' \vdash K \approx_r^{\mathcal{P}} K' :^{\mathcal{P}} \Lambda; \Lambda'$ and $L \approx_r^{\mathcal{P}} L' :^{\mathcal{P}} \Lambda'$ then $\text{Bdy}_{L, L'}^{\mathcal{P}} K K' (\Lambda; \Lambda') \subseteq \approx_r^{\mathcal{P}}$.

Proof. By coinduction via Corollary 1. We have to show $\text{Bdy}_{L, L'}^{\mathcal{P}} F F' (\Lambda; \Lambda') \subseteq \lesssim_{r'}$ from $r \subseteq r'$ and the coinductive hypothesis $\text{Bdy}_{L, L'}^{\mathcal{P}} F F' (\Lambda; \Lambda') \subseteq r'$. Applying clause (3) of the first premise reduces the proof obligation to $K; L \approx_{r'}^{\mathcal{P}} K; L' :^{\mathcal{P}} \Lambda; \Lambda'$. We apply Lemma 13. The third premise of Lemma 13 follows from the coinductive hypothesis and $r' \subseteq \lesssim_{r'} \cup r'$, the fourth premise follows from monotonicity (Definition 6 and Lemma 11). \square

Lemma 14 shows that equivalence of function bodies is sufficient to show that the corresponding recursive functions are equivalent.

Lemma 15 (Extension) Let \mathcal{P} be a proof relation. If we have $L \mid L' \vdash K \approx_r^{\mathcal{P}} K' :^{\mathcal{P}} \Lambda; \Lambda'$ and $L \approx_r^{\mathcal{P}} L' :^{\mathcal{P}} \Lambda'$ then $K; L \approx_r^{\mathcal{P}} K; L' :^{\mathcal{P}} \Lambda; \Lambda'$.

Proof. We apply Lemma 13. The only non-trivial premise is to show

$$\text{Bdy}_{L,L'}^{\mathcal{P}} K K' (\Lambda; \Lambda') \subseteq \approx_r^{\mathcal{P}} \cup r$$

We make use of the fact $\approx_r^{\mathcal{P}} \subseteq \approx_r^{\mathcal{P}} \cup r$ and finish the proof with Lemma 14. \square

Lemma 16 (Fun Compatibility) Let \mathcal{P} be a proof relation. If $L \mid L' \vdash \langle F \rangle_V \approx_r^{\mathcal{P}} \langle F' \rangle_{V'} :^{\mathcal{P}} \Lambda; \Lambda'$ and $L \approx_r^{\mathcal{P}} L' :^{\mathcal{P}} \Lambda; \Lambda'$ and

$$\begin{aligned} \forall r, \langle F \rangle_V; L \approx_r^{\mathcal{P}} \langle F' \rangle_{V'}; L' :^{\mathcal{P}} \Lambda; \Lambda' \rightarrow \\ \langle F \rangle_V; L, V, t \approx_r^{\mathcal{P}} \langle F' \rangle_{V'}; L', V', t' \end{aligned}$$

then $(L, V, \text{fun } F \text{ in } t) \approx_r^{\mathcal{P}} (L', V', \text{fun } F' \text{ in } t')$.

Proof. We reduce both sides one step. We apply the last premise and have to show $\langle F \rangle_V; L \approx_r^{\mathcal{P}} \langle F' \rangle_{V'}; L' :^{\mathcal{P}} \Lambda; \Lambda'$. Lemma 15 finishes the proof. \square

When using Lemma 15 or Lemma 16 it suffices to show that the bodies of new function definitions are related according to Definition 15. Item (3) already provides that the contexts containing the new functions are related. See the function definition case of Lemma 18 to understand how this enables the inductive method.

6.3 Using Related Function Contexts to Prove the Application Case

Definition 16 Argument evaluation of L, V, \bar{e} and L', V', \bar{e}' agrees with respect to p and \mathcal{P} if whenever $L_f = (V, \bar{x}, s)$ and $L_{f'} = (V, \bar{x}', s')$ and $\text{Param } a \bar{x} \bar{x}'$ then

1. if $\llbracket \bar{e} \rrbracket V = \bar{v}$ and $|\bar{x}| = |\bar{v}|$ then there exists \bar{v}' such that $\llbracket \bar{e}' \rrbracket V = \bar{v}'$ and $|\bar{x}'| = |\bar{v}'|$ and $\text{Arg } a \bar{v} \bar{v}'$
2. if $p = \text{bisim}$ and $\llbracket \bar{e} \rrbracket V = \perp$ then $\llbracket \bar{e}' \rrbracket V' = \perp$
3. if $p = \text{bisim}$ and $|\bar{x}| \neq |\bar{e}|$ then $|\bar{x}'| \neq |\bar{e}'|$

Lemma 17 Let \mathcal{P} be a proof relation. If

1. $L \approx_r^{\mathcal{P}} L' :^{\mathcal{P}} \Lambda$
2. $\text{Idx } \Lambda_f f f'$
3. argument evaluation of L, V, \bar{e} and L', V', \bar{e}' agrees with respect to p and \mathcal{P}

then $(L, V, f \bar{e}) \approx_r^{\mathcal{P}} (L', V', f' \bar{e}')$.

Proof. From (2) we have that Λ_f is defined, and by definition of (1) $\text{dom } \Lambda = \text{dom } L$, which means $f \in \text{dom } L$, hence again by definition of (2) $f' \in \text{dom } L'$. We assume that $L_f = (V, \bar{x}, s)$ and $L_{f'} = (V, \bar{x}', s')$. By definition of (1) we know $\text{Param } a \bar{x} \bar{x}'$, hence $\text{Param } \Lambda_f \bar{x} \bar{x}'$. Case analysis.

- Case $\llbracket \bar{e} \rrbracket V = \bar{v}$.
 - If $|\bar{x}| = |\bar{e}|$ we exploit clause (1) of premise (3) and obtain the fact

$$\text{App}^{\mathcal{P}} \wedge L L' (L, V, f \bar{e}) (L', V', f' \bar{e}')$$

We know $\text{App}^{\mathcal{P}} \wedge L L' \subseteq r$ from premise (1) and are done.

- If $|\bar{x}| \neq |\bar{e}|$. If $p = \text{sim}$, we are done using SIM-ERROR. If $p = \text{bisim}$, we exploit clause (3) of assumption (3) and obtain that $|\bar{x}'| \neq |\bar{e}'|$. Both sides are stuck (SIM-TERM).
- Case $\llbracket \bar{e} \rrbracket V = \perp$. If $p = \text{sim}$, we are done by SIM-ERROR. If $p = \text{bisim}$, we exploit clause (2) of assumption (3) and obtain that $\llbracket \bar{e}' \rrbracket V' = \perp$. Both sides are stuck. \square

7 Bisimilarity and Similarity are Contextual

In this section we use the inductive method to show that bisimilarity and similarity are contextual, that is, sound for contextual equivalence and contextual approximation.

Definition 17 We define the proof relation \mathcal{P}_{eq} where

$$\begin{aligned} A &:= \text{list } \mathcal{V} \\ \text{Param } \bar{x} \bar{y} \bar{y}' &:= \bar{x} = \bar{y} \wedge \bar{y} = \bar{y}' \\ \text{Arg } \bar{x} \bar{v} \bar{v}' &:= V = V' \wedge \bar{v} = \bar{v}' \wedge |\bar{x}| = |\bar{v}| \\ \text{Idx } _ f f' &:= f = f' \end{aligned}$$

Let $\boxed{\text{pa } L}$ denote the projection of each closure to its parameters, and $\boxed{\text{pa } F}$ the projection of each function definition to its parameters.

Definition 18 (Program Equivalence) For two terms s, s' we define equivalence

$$s \simeq_r^{\mathcal{P}} s' \text{ as } \boxed{s \simeq_r^{\mathcal{P}} s'}$$

$$\forall LL'V, L \approx_r^{\mathcal{P}} L' :^{\mathcal{P}_{\text{eq}}} (\text{pa } L') \rightarrow (L, V, s) \approx_r^{\mathcal{P}} (L', V, s')$$

Lemma 18 (Reflexivity) $s \simeq_r^{\mathcal{P}} s$.

Proof. By induction on s .

- The case for let follows from the inductive hypothesis, and lemmas SIM-LET-OP and SIM-LET-CALL.
- The conditional case follows by SIM-COND and the inductive hypotheses.
- In the case of application $f \bar{e}$, we do a case analysis on whether L'_f exists. If it exists we are done by $L \approx_r^{\mathcal{P}} L' :^{\mathcal{P}_{\text{eq}}} \text{pa } L'$ with Lemma 17, and the fact that argument evaluation obviously agrees. Otherwise, we know from $L \approx_r^{\mathcal{P}} L' :^{\mathcal{P}_{\text{eq}}} \text{pa } L'$ that $\text{dom } L = \text{dom } L'$, so L_f does not exist either, and both sides are stuck. We finish with SIM-TERM.

- The case for operation is by SIM-TERM.
- In the function definition case we apply Lemma 16. The second premise of Lemma 16 is an assumption, and the third is the inductive hypothesis. For the first premise $L \mid L' \vdash \langle F \rangle_V \approx_r^p \langle F' \rangle_{V'} :^P \Lambda; \Lambda'$ (Definition 15) we need to show separation and parameter relation, which hold because both sides define the same functions. It remains to show (3) of Definition 15, i.e. from

$$(*) \langle F \rangle_V; L \approx_r^p \langle F \rangle_V; L' :^P (\text{pa } F; \text{pa } L')$$

that $\text{Bdy}_{L,L'}^P \langle F \rangle_V \langle F \rangle_V (\text{pa } F; \text{pa } L') \subseteq \approx_r^p$ holds. Unfolding Bdy (Definition 13), we get $\text{Idx} (\text{pa } F; \text{pa } L') f f'$ and $F_f = (\bar{x}, s)$ and $F_{f'} = (\bar{x}', s')$ and $\text{Arg } \bar{x} \bar{v} \bar{v}'$. After unfolding Idx to obtain $f = f'$, and after unfolding Arg to obtain $\bar{v} = \bar{v}'$ we have to show that

$$(\langle F \rangle_V; L, V[\bar{x} \mapsto v], s) \sim_r (\langle F \rangle_V; L', V[\bar{x} \mapsto v], s)$$

This follows from the inductive hypothesis, but only with (*). Note that if Lemma 16 had not provided (*) through Definition 15, the proof would not work. \square

Lemma 19 (Reflexivity) $L \approx_r^p L :^P \text{pa } L$.

Lemma 20 (Transitivity) $s \simeq_\perp^p s' \rightarrow s' \simeq_\perp^p s'' \rightarrow s \simeq_\perp^p s''$.

Theorem 1 Let $C[\]$ be an IL context, i.e. a term with a hole. $(\forall r, s \simeq_r^p s') \rightarrow C[s] \simeq_r^p C[s']$.

Proof. Induction on C . \square

Lemma 21 If we have $\forall r, i, s_i \simeq_r^p s'_i$ and $\forall r, t \simeq_r^p t'$ then $\forall r, \text{fun } \overline{f \bar{x} = s} \text{ in } t \simeq_r^p \text{fun } \overline{f \bar{x} = s'} \text{ in } t'$.

8 Unreachable Code Elimination

We apply our inductive proof method in the correctness proof of the first optimization: unreachable code elimination.

8.1 Representing Program Analysis Information

A program analysis associates information with every subterm of a program. From now on, we use annotated terms instead of terms, i.e., $\{a\} s$, which associates analysis information a with term s . Note that the subterms in s are again annotated. Given analysis information of type A , we denote this new inductive type by $\text{Ann Exp } A$. The formal development contains verified analyses for DVE and UCE.

$$\begin{array}{c}
\text{REACH-LET} \quad \frac{b = b' \quad \Lambda \vdash \mathbf{reach} \{b'\} s}{\Lambda \vdash \mathbf{reach} \{b\} \text{let } x = \eta \text{ in } \{b'\} s} \quad \text{LIVE-EXP} \quad \frac{}{\Lambda \vdash \mathbf{reach} \{b\} e} \quad \text{REACH-APP} \quad \frac{b \rightarrow \Lambda_f}{\Lambda \vdash \mathbf{reach} \{b\} f \bar{e}} \\
\\
\text{REACH-COND} \quad \frac{\beta(\llbracket e \rrbracket \emptyset) \neq \mathbf{false} \rightarrow b = b_1 \quad \Lambda \vdash \mathbf{reach} \{b_1\} s_1 \quad \beta(\llbracket e \rrbracket \emptyset) \neq \mathbf{true} \rightarrow b = b_2 \quad \Lambda \vdash \mathbf{reach} \{b_2\} s_2}{\Lambda \vdash \mathbf{reach} \{b\} \text{if } e \text{ then } \{b_1\} s_1 \text{ else } \{b_2\} s_2} \\
\\
\text{REACH-FUN} \quad \frac{\forall g, \overline{f : b}; \Lambda \vdash \mathbf{reach} \{b_g\} s_g \quad \overline{f : b}; \Lambda \vdash \mathbf{reach} \{c\} t \quad a = c}{\Lambda \vdash \mathbf{reach} \{a\} \text{fun } f \bar{x} = \{b\} s \text{ in } \{c\} t}
\end{array}$$

Figure 8: Definition of the Reachability Predicate $\Lambda \vdash \mathbf{reach} s$. The context $\Lambda : \text{context } \mathbb{B}$ contains reachability information for functions and $s : \text{Ann Exp } \mathbb{B}$ is an program annotated with reachability information.

8.2 Inductive Reachability Judgment

We want to annotate t with reachability information $b, b' : \mathbb{B}$ such that whenever $(L, V, \{b\} s) \rightarrow^* (L', V', \{b'\} t)$ we have $b \rightarrow b'$. Reachability is a non-trivial semantic property, hence undecidable. We define inductively the judgment **reach** in Figure 8. Reach considers the value of the condition expression e in the empty environment: $\llbracket e \rrbracket \emptyset$. If $\llbracket e \rrbracket \emptyset = \perp$, both cases are assumed to be reachable, which may over-approximate. REACH-LET propagates reachability information through let-bindings: If the let is reachable, then so is its successor. REACH-APP ensures that whenever a function application is reachable, then the function is also reachable. REACH-COND evaluates $\llbracket e \rrbracket \emptyset$, i.e. it evaluates the condition under the empty variable environment. If $\beta(\llbracket e \rrbracket \emptyset) = \mathbf{false}$ then reachability is not propagated into the consequence s_1 . Propagation into the alternative s_2 is treated similarly. REACH-FUN propagates reachability into t . The context Λ is extended with the reachability information of the function bodies $\overline{f : b}$. The topmost premise ensures that the reachability information for all function bodies is sound.

8.3 Transformation and Correctness

In Figure 10, we define a function **uce** that removes all code not marked reachable. If all functions from a mutually recursive function definition are removed, the **fun**-statement is removed, too. Conditionals are removed if the value of the condition can be statically evaluated.

```

filterby :  $\forall XY, (X \rightarrow \mathbb{B}) \rightarrow \text{list } X \rightarrow \text{list } Y \rightarrow \text{list } Y$ 
filterby  $p (x :: x') (y :: y') =$ 
  if  $px$  then  $y :: \text{filterby } p x' y'$  else  $\text{filterby } p x' y'$ 
filterby  $p \_, \_ = \text{nil}$ 
filter :  $\forall X, (X \rightarrow \mathbb{B}) \rightarrow \text{list } X \rightarrow \text{list } X$ 
filter  $p x = \text{filterby } p x x$ 

```

Figure 9: Definition of filter

```

uce :  $\text{Ann Exp } \mathbb{B} \rightarrow \text{Exp}$ 
uce (let  $x = \eta$  in  $s$ ) = let  $x = \eta$  in (uce  $s$ )
uce (if  $e$  then  $s_1$  else  $s_2$ ) =
  if  $\llbracket e \rrbracket \emptyset = \text{true}$  then uce  $s_1$ 
  else if  $\llbracket e \rrbracket \emptyset = \text{false}$  then uce  $s_2$ 
  else if  $e$  then (uce  $s_1$ ) else (uce  $s_2$ )
uce ( $f \bar{e}$ ) =  $f \bar{e}$ 
uce  $e = e$ 
uce (fun  $F$  in  $t$ ) =
  let  $F' = \text{uceF } F$  in
  if  $|F'| = 0$  then uce  $t$  else fun  $F'$  in (uce  $t$ )
uceF  $F =$ 
  let  $K = \text{filter } (\lambda(\bar{x}, \{b\} s). b) F$  in
  map ( $\lambda(\bar{x}, s). (\bar{x}, \text{uce } s)$ )  $K$ 

```

Figure 10: Definition of Unreachable Code Elimination

Definition 19 We define the proof relation \mathcal{P}_{uce} where

$$\begin{aligned} A &:= \mathbb{B} \\ \text{Param } _ \bar{x} \bar{x}' &:= \bar{x} = \bar{x}' \\ \text{Arg } a \bar{v} \bar{v}' &:= \bar{v} = \bar{v}' \\ \text{Idx } a f f' &:= a = \mathbf{true} \wedge f = f' \end{aligned}$$

Lemma 22 If $\text{dom } \Lambda = \text{dom } F$ then $F \parallel \Lambda; \Lambda' \parallel_{\mathcal{P}_{\text{uce}}} \text{uceF } F$.

Lemma 23 $F = \overline{f \bar{x} = \{b\} s} \rightarrow \text{Param } \bar{b} F (\text{uceF } F)$.

Lemma 24 Suppose that it holds $F' = \text{uceF } F$ and we have that $(\llbracket F \rrbracket_V; L, V, s) \sim_r (\llbracket F' \rrbracket_{V'}; L', V', \text{uce } t)$ then

$$(L, V, \text{fun } F \text{ in } s) \sim_r (L', V', \text{if } |F'| = 0 \text{ then } \text{uce } t \text{ else fun } F' \text{ in } \text{uce } t).$$

Proof. If $|F'| = 0$, then $F' = \text{nil}$. After reducing only the left side one step via SIM-EXPANSION-CLOSED, the assumption solves the goal. Otherwise we reduce both sides one step (Lemma 2), and the assumption solves the goal. \square

Theorem 2 Let $\Lambda \vdash \mathbf{reach} \{\mathbf{true}\} s$ and $L \sim_r L' :^{\mathcal{P}_{\text{uce}}} \Lambda$. Then: $(L, V, s) \sim_r (L', V, \text{uce } s)$.

Proof. Induction on s and in each case inversion of **reach**.

- The case for let follows from SIM-LET-CALL and the inductive hypothesis.
- The case for the conditional follows by Lemma 10 and the inductive hypotheses.
- The application case follows from $L \sim_r L' :^{\mathcal{P}_{\text{uce}}} \Lambda$ with Lemma 17, after discharging premises: Note that $\Lambda_f = \mathbf{true}$ by inversion on **reach**, so $\text{Idx } \Lambda_f f f'$ holds by definition. Argument evaluation agrees, since parameters, and arguments and environments are identical.
- The case for operation is trivial, since operation and environments are identical.
- In the function definition case, let $F' = \text{uceF } F$. Lemma 24 lets us deal with both cases uniformly and requires

$$(\llbracket F \rrbracket_V; L, V, t) \sim_r (\llbracket F' \rrbracket_E; L', E, \text{uce } t)$$

After applying the inductive hypothesis, we must show $\llbracket F \rrbracket_E; L \lesssim_r \llbracket F' \rrbracket_E; L' :^{\mathcal{P}} \Lambda'; \Lambda$. We apply Lemma 15 and discharge its premises by using Lemma 22 and Lemma 23. The remaining premise requires us to show from

$$\llbracket F \rrbracket_E; L \sim_r \llbracket F' \rrbracket_E; L' :^{\mathcal{P}} \Lambda'; \Lambda \quad (*)$$

that $\text{Bdy}_{L, L'}^{\mathcal{P}} (\llbracket F \rrbracket_E (\llbracket F' \rrbracket_E (\Lambda'; \Lambda)) \subseteq \sim_r$. Unfolding Bdy, we obtain $\text{Idx } (\Lambda'; \Lambda) f f'$ and $F_f = (\bar{x}, s)$ and $F'_{f'} = (\bar{x}', s')$ and $\text{Arg } \Lambda'_f \bar{v} \bar{v}'$. Unfolding those, we have to show

$$(\llbracket F \rrbracket_E; L, E[\bar{x} \mapsto v], s) \sim_r (\llbracket F' \rrbracket_E; L', E[\bar{x} \mapsto v], \text{uce } s)$$

The inductive hypothesis solves the goal with $(*)$. \square

9 Dead Variable Elimination

Dead variable elimination relies on a true liveness analysis. A variable is live if it is (potentially) used later on. A variable is true live, if it is used to compute a value that is live later on. True liveness requires a fixed-point computation, but is able to detect parameters of a function that do not contribute to the behavior of the function.

9.1 Inductive Liveness Judgement

We specify sound true liveness information with the inductive judgment **tlive** in Figure 11.

TLIVE-OP ensures that all variables that are live after the let (X') are also live before the let, except the variable defined. The free variables of e only need to be live if x is live after the let. TLIVE-CALL is similar, but always requires the free variables of \bar{e} to be live, as we can never remove calls (even if their result is unused). TLIVE-EXP requires the free variables of \bar{e} to be live. TLIVE-COND tests whether the condition is a constant expression by evaluating it in the empty environment. Only if this is unsuccessful, we require its free variables to be live. If $\beta(\llbracket e \rrbracket \emptyset) \neq \mathbf{false}$, the consequence might be reachable. In this case the rule requires that all variables live in the consequence (X_1) are live before the conditional ($X_1 \subseteq X$), and that the judgment holds recursively. Otherwise, no requirements are imposed. TLIVE-APP requires whenever a parameter x_i is in the live set of the function body X_f , then the free variables of the corresponding argument expression e_i are live at the application. TLIVE-FUN requires the liveness judgment to recursively hold for the continuation and the function bodies under extended contexts. The context ζ is extended with the parameters $\bar{f} : \bar{x}$ of the newly defined functions, and the context Λ is extended with the liveness information $\bar{f} : \bar{Y}$ for the function bodies. The rule requires all variables live after the function definition to be live before the function definition, and that all variables live in the function bodies (except parameters) are also live before the function definition.

9.2 Transformation and Correctness

We realize dead variable elimination (DVE) with the recursive function **dve** defined in Figure 12. The recursive procedure descends through the program, removes unused let-bindings and filters parameter and argument lists according to liveness information. As in UCE, conditionals are removed if the condition can be statically evaluated.

Definition 20 We define the proof relation \mathcal{P}_{dve} where

$$\begin{aligned} A &:= \text{list } \mathcal{V} \times \text{set } \mathcal{V} \\ \text{Param } (\bar{x}, X) \bar{y} \bar{y}' &:= \bar{x} = \bar{y} \wedge \bar{y}' = \text{filter}(\lambda x. x \in X) \bar{y} \\ \text{Arg } (\bar{x}, X) \bar{v} \bar{v}' &:= \bar{v}' = \text{filterby}(\lambda(x). x \in X) \bar{x} \bar{v} \\ &\quad \wedge |\bar{x}| = |\bar{v}| \\ \text{Idx } _ f f' &:= f = f' \end{aligned}$$

$$\begin{array}{c}
\text{TLIVE-OP} \\
\frac{X' \setminus \{x\} \subseteq X \quad x \in X' \rightarrow \text{fv}(e) \subseteq X \quad \zeta \mid \Lambda \vdash \mathbf{tlive} \{X'\} s}{\zeta \mid \Lambda \vdash \mathbf{tlive} \{X\} \text{ let } x = e \text{ in } \{X'\} s} \quad \text{TLIVE-EXP} \frac{\text{fv}(e) \subseteq X}{\zeta \mid \Lambda \vdash \mathbf{tlive} \{X\} e} \\
\\
\text{TLIVE-CALL} \frac{X' \setminus \{x\} \subseteq X \quad \text{fv}(\bar{e}) \subseteq X \quad \zeta \mid \Lambda \vdash \mathbf{tlive} \{X'\} s}{\zeta \mid \Lambda \vdash \mathbf{tlive} \{X\} \text{ let } x = \alpha \bar{e} \text{ in } \{X'\} s} \quad \text{TLIVE-APP} \frac{|\bar{x}| = |\bar{e}| \quad \forall i, x_i \in \Lambda_f \rightarrow \text{fv}(e_i) \subseteq X}{\zeta \mid \Lambda \vdash \mathbf{tlive} \{X\} f \bar{e}} \\
\\
\text{TLIVE-COND} \frac{\llbracket e \rrbracket \emptyset = \perp \rightarrow \text{fv}(e) \subseteq X \quad \beta(\llbracket e \rrbracket \emptyset) \neq \mathbf{false} \rightarrow X_1 \subseteq X \wedge \zeta \mid \Lambda \vdash \mathbf{tlive} \{X_1\} s_1 \quad \beta(\llbracket e \rrbracket \emptyset) \neq \mathbf{true} \rightarrow X_2 \subseteq X \wedge \zeta \mid \Lambda \vdash \mathbf{tlive} \{X_2\} s_2}{\zeta \mid \Lambda \vdash \mathbf{tlive} \{X\} \text{ if } e \text{ then } \{X_1\} s_1 \text{ else } \{X_2\} s_2} \\
\\
\text{TLIVE-FUN} \frac{\overline{f : \bar{x}}; \zeta \mid \overline{f : Y}; \Lambda \vdash \mathbf{tlive} \{X'\} t \quad X' \subseteq X \quad \forall g, \overline{f : \bar{x}}; \zeta \mid \overline{f : Y}; \Lambda \vdash \mathbf{tlive} \{Y_g\} s_g \quad \forall g, Y_g \setminus \bar{x}_g \subseteq X}{\zeta \mid \Lambda \vdash \mathbf{tlive} \{X\} \text{ fun } \overline{f \bar{x} = \{Y\}} s \text{ in } \{X'\} t}
\end{array}$$

Figure 11: Definition of the True Liveness Predicate $\zeta \mid \Lambda \vdash \mathbf{tlive} s$. The context $\zeta : \text{context}(\text{list } \mathcal{V})$ contains parameters of functions, $\Lambda : \text{context}(\text{set } \mathcal{V})$ contains the variables live in the function body, and $s : \text{Ann Exp}(\text{set } \mathcal{V})$ is a program annotated with liveness information.

$$\begin{aligned}
& \text{dve} : \text{list } \mathcal{V} \rightarrow \text{set } \mathcal{V} \rightarrow \text{Ann Exp (set } \mathcal{V}) \rightarrow \text{Exp} \\
& \text{dve } \zeta \Lambda (\text{let } x = e \text{ in } \{X\} s) = \\
& \quad \text{let } s' = \text{dve } \zeta \Lambda (\{X\} s) \text{ in} \\
& \quad \text{if } x \in X \text{ then let } x = e \text{ in } s' \text{ else } s' \\
& \text{dve } \zeta \Lambda (\text{let } x = \alpha \bar{e} \text{ in } s) = \\
& \quad \text{let } x = \alpha \bar{e} \text{ in } (\text{dve } \zeta \Lambda s) \\
& \text{dve } \zeta \Lambda (\text{if } e \text{ then } s_1 \text{ else } s_2) = \\
& \quad \text{if } \llbracket e \rrbracket \emptyset = \mathbf{true} \text{ then } \text{dve } \zeta \Lambda s_1 \\
& \quad \text{else if } \llbracket e \rrbracket \emptyset = \mathbf{false} \text{ then } \text{dve } \zeta \Lambda s_2 \\
& \quad \text{else if } e \text{ then } (\text{dve } \zeta \Lambda s_1) \text{ else } (\text{dve } \zeta \Lambda s_2) \\
& \text{dve } (\zeta; f : \bar{x}; \zeta') (\Lambda; f : X; \Lambda') (\{_ \} f \bar{e}) = \\
& \quad f (\text{filterby } (\lambda x. x \in X) \bar{x} \bar{e}) \\
& \text{dve } \zeta \Lambda e = e \\
& \text{dve } \zeta \Lambda (\text{fun } \overline{f \bar{x}} = \{X\} s \text{ in } t) = \\
& \quad \text{let } \forall i, F'_i = (\text{filter } (\lambda x. x \in X_i) \bar{x}_i, \\
& \quad \quad \text{dve } (\bar{x}; \zeta) (\bar{X}; \Lambda) s_i) \text{ in} \\
& \quad \text{fun } F' \text{ in } (\text{dve } (\bar{x}, \zeta) (\bar{X}; \Lambda) t)
\end{aligned}$$

Figure 12: Definition of Dead Variable Elimination

Lemma 25 If $\text{dom } F = \text{dom } \Lambda = \text{dom } F'$ then we have $F \parallel \Lambda; \Lambda' \parallel_{\mathcal{P}_{\text{dve}}} F'$.

Lemma 26 Let $F = \overline{f \bar{x} = \{X\} s}$ and F' s.t. for all i ,

$$F'_i = (\text{filter } (\lambda x. x \in X_i) \bar{x}_i, \text{dve } (\bar{\bar{x}}; \zeta) (\bar{X}; \Lambda) s_i)$$

and $|F| = |F'|$. Then $\text{Param } (\overline{f : (\bar{x}, X)}) F F'$.

$$V =_X V'$$

We are now ready to show the correctness theorem. We write $\boxed{V =_X V'}$ if V and V' agree on the values of the variables in the set X .

Theorem 3 Let $\zeta \mid \Lambda \vdash \mathbf{tlive} \{X\} s$ and $L \lesssim_r L' :^{\mathcal{P}_{\text{dve}}} \text{zip } \zeta \Lambda$ and $V =_X V'$. Then:

$$(L, V, s) \lesssim_r (L', V', \text{dve } \zeta \Lambda \{X\} s).$$

Proof. Induction on s and in each case inversion of **tlive**.

- The case for let-call follows from SIM-LET-CALL and the inductive hypothesis.
- In the case for let op, we do a case analysis in $x \in X$.
 - If $x \in X$, the case follows from SIM-LET-OP with the fact that $\llbracket e \rrbracket V = \llbracket e \rrbracket V'$ because V and V' agree on $\text{fv}(e)$.
 - If $x \notin X$, case analysis on $\llbracket e \rrbracket V$.
 - * If $\llbracket e \rrbracket V = \perp$, the left side is stuck (SIM-ERROR).
 - * If $\llbracket e \rrbracket V = v$, we use SIM-EXPANSION-CLOSED to reduce the left side one step and are done by the inductive hypothesis with the observation that $V[x \mapsto v]$ still agrees with V' on the X because $x \notin X$.
- The case for the conditional follows by Lemma 10 and the inductive hypotheses.
- The case for application follows from $L \sim_r L' :^{\mathcal{P}_{\text{dve}}} \text{zip } \zeta \Lambda$ with Lemma 17, after discharging premises. The relation $\text{Idx } (\bar{x}, X) f f$ holds by definition. Argument evaluation agrees because

$$\begin{aligned} & \llbracket \text{filterby } (\lambda x. x \in X_f) \bar{x} \bar{e} \rrbracket V' \\ &= \text{filterby } (\lambda x. x \in X_f) \bar{x} (\llbracket \bar{e} \rrbracket V) \end{aligned}$$

since we already know that $\llbracket \bar{e} \rrbracket V \neq \perp$ and V and V' agree on the live variables.

- The case for operation is trivial, since operation are identical and environments agree on the live variables.
- In the function definition case, we have F' such that $|F'| = |F|$ and for all i

$$F'_i = (\text{filter } (\lambda x. x \in X_i) \bar{x}_i, \text{dve } (\bar{\bar{x}}; \zeta) (\bar{X}; \Lambda) s_i)$$

We apply Lemma 16 and have to discharge premises. The second premise holds by assumption, the third is the inductive hypothesis. The first two requirements of the

first premise are Lemma 25 and Lemma 26. It remains to show from $(*) \llbracket F \rrbracket; L \sim_r \llbracket F' \rrbracket; L' :^P (\overline{x}, \overline{X}; \Lambda)$ that

$$\text{Bdy}_{L, L'}^P F F' (\overline{x}, \overline{X}; \Lambda) \subseteq \sim_r$$

After unfolding Bdy (Definition 13), we have the assumptions $\text{Idx}(\overline{x}, X) f f'$ and $F_f = (\overline{x}, s)$ and $F_{f'} = (\overline{x}', s')$ and $\text{Arg}(\overline{x}, X) \overline{v} \overline{v}'$. And after further unfolding we get $\overline{x}' = \text{filter}(\lambda x. x \in X) \overline{x}$ and $\overline{v}' = \text{filterby}(\lambda(x). x \in X) \overline{x} \overline{v}$. We have to show that

$$\begin{aligned} & (\llbracket F \rrbracket; L, V[\overline{x} \mapsto v], s) \\ \sim_r & (\llbracket F' \rrbracket; L', V'[\overline{x}' \mapsto \overline{v}'], \text{dve}(\overline{\overline{x}}; \zeta)(\overline{X}; \Lambda) s) \end{aligned}$$

Inductive hypothesis provides the latter. Its premises are discharged by $(*)$ and the observation that the updated environments still agree on the live variables. \square

10 Coq Development

The formal development accompanying this paper is part of a verified compiler LVC. LVC use the inductive method presented in this paper, and variations of it, for many correctness proofs. These include DVE, UCE, Copy Propagation, Sparse Conditional Constant Propagation, and some lowering passes. LVC also features an imperative variant of IL, which is called IL/I and serves as source language. The difference between IL and IL/I is that the latter uses imperative variables instead of lexically scoped binders. The first transformations in the pipeline of the LVC compiler are UCE and DVE on IL/I. The formal development hence also contains proofs of UCE and DVE for IL/I. The setup as described works for IL/I, too, and the proof structure remains the same.

In the formal development we use De-Bruijn indices, not a named representation for function binders. This fact complicates UCE, as indices change whenever a function is removed. Fortunately, this does not causes problems with our inductive method.

The part of the LVC development that pertains to this paper is available online:

www.ps.uni-saarland.de/~sdschn/lvc-ind/

LVC has more than 36k LoC. The formalization of the inductive method (Section 6) takes ~ 400 LoC. The correctness theorems (Theorem 2, Theorem 3) are ~ 50 LoC each. Also counting lemmas, it takes ~ 300 LoC to show each of DVE and UCE correct.

11 Conclusion

We described an inductive method for proofs of simulation-based program equivalence. In contrast to the standard approach, which indexes the (bi)simulation with a measure, our approach works without modifying the simulation. With our method, bisimilarity can be proven with the need for symmetrization.

After the method is setup, the overhead of the correctness proofs of transformations is low, and the proof becomes a straight-forward induction. The details of the proof are simple enough to be explained in full on paper. The method separates concerns: The correctness proofs follow the syntactic definition of the transformation, and a separate, general lemma proved by coinduction is used to deal with fixed-point computation in the language. This allows to focus on the actual verification problem inherent to the transformation instead of requirements induced by the proof method. We think inductive methods like ours are an essential tool for bisimulation-based compiler verification and useful in general.

We applied the method to two optimizations, unreachable code elimination (UCE) and dead variable elimination (DVE). The optimizations are not straight-forward to verify because they remove instructions and change function signatures. The inductive method also improves modularity of the correctness argument: We argued certain removal steps in a separate lemma (e.g. Lemma 10). After using such a lemma in a plain coinductive proof (even when using Paco [3]), further justification would be required before the cohypothesis could be applied.

References

- [1] Gilles Barthe, Delphine Demange, and David Pichardie. “A Formally Verified SSA-Based Middle-End - Static Single Assignment Meets CompCert”. In: *ESOP*. Vol. 7211. LNCS. Tallinn, Estonia, Mar. 24–Apr. 1, 2012.
- [2] Douglas J. Howe. “Proving Congruence of Bisimulation in Functional Programming Languages”. In: *Inf. Comput.* 124.2 (1996).
- [3] Chung-Kil Hur, Georg Neis, Derek Dreyer, and Viktor Vafeiadis. “The power of parameterization in coinductive proof”. In: *POPL*. Rome, Italy, Jan. 23–25, 2013.
- [4] Xavier Leroy. “A Formally Verified Compiler Back-end”. In: *JAR* 43.4 (2009).
- [5] Xavier Leroy. “Formal Verification of a Realistic Compiler”. In: *CACM* 52.7 (2009).
- [6] Robin Milner, Mads Tofte, and David Macqueen. *The Definition of Standard ML*. Cambridge, MA, USA: MIT Press, 1997.
- [7] Georg Neis et al. “Pilsner: a compositionally verified compiler for a higher-order imperative language”. In: *ICFP*. Vancouver, Canada, Sept. 1–3, 2015.
- [8] Scott Owens, Magnus O. Myreen, Ramana Kumar, and Yong Kiam Tan. “Functional Big-Step Semantics”. In: *ESOP*. Vol. 9632. LNCS. Eindhoven, The Netherlands, May 2–8, 2016.
- [9] Sigurd Schneider. “Semantics of an Intermediate Language for Program Transformation”. Master’s Thesis. Saarland University, 2013. URL: <http://www.ps.uni-saarland.de/~sdschn/master>.
- [10] Sigurd Schneider, Gert Smolka, and Sebastian Hack. “A Linear First-Order Functional Intermediate Language for Verified Compilers”. In: *ITP*. Vol. 9236. LNCS. Nanjing, China, Aug. 24–27, 2015.
- [11] Jaroslav Sevcík, Viktor Vafeiadis, Francesco Zappa Nardelli, Suresh Jagannathan, and Peter Sewell. “CompCertTSO: A Verified Compiler for Relaxed-Memory Concurrency”. In: *J. ACM* 60.3 (2013).

- [12] Jianzhou Zhao, Santosh Nagarakatte, Milo M. K. Martin, and Steve Zdancewic. “Formal Verification of SSA-based Optimizations for LLVM”. In: *PLDI*. Seattle, WA, USA, June 16–19, 2013.